

SOS10 Vendor Panel: SGI Directions for Productive, Profitable PFLOPS Computing


Steve Reinhardt
Principal Engineer
spr <at> sgi <dot> com






SGI is shifting the focus of its system architecture, to accelerating the discovery of knowledge in immense data.

Scientific Process



**Who coined this???

Pattern Discovery in Biology**: Metabolic Behavior of Communities of Cells



**Other examples from climate change research


System Requirements for Knowledge Discovery

1. Major commodity content for feasibility/affordability
2. **Big memory (PB by 2009)**
3. Big data ingest/storage (150+Gb/s/link, 80GB/s to a single FS, 100PB FS)
4. Heterogeneous processing
 - FPGAs, special-purpose processors (e.g., ClearSpeed, Cell)
5. RAS for production-worthiness
6. **Rapid application development**
 - Star-P, Graph Analysis and Pattern Discovery toolbox
7. In-the-loop visualization of big data (1TB)

Big Memory

- Need to consider the whole graph at once -> PB
- Graphs typically don't exhibit locality -> GAM, global strided/random, large-span TLBs
- Graph algorithms don't use simple synchronization -> irregular, data-driven synch among dynamic groups of processors
- Info is crucially valuable -> RAS

Rapid Application Development



```
a = rand(n,n*p);  
ppload imagedata a  
[nrow ncol] = size(a);  
b = ones(nrow,ncol);  
c = fft2(a);  
d = ifft2(c);  
  
diff = max(max(abs(a-d)));  
if (diff > 10*eps)  
    sprintf('Error, diff=%f',  
diff);  
end
```

• Star-P

- extends MATLAB language for ||ism
- very high (>95%) reuse of existing MATLAB code
- retains interactivity
- addresses new audience of users
- builds on parallel libraries
- demonstrated 2.8TB array
- causes MPI usage to plummet

• Pattern Discovery

- cSSCA#2:
 - 110 SLOC total, 5X fewer than MATLABmpi, 9X fewer than C/pthreads
 - strong scaling to 10^9 edges/128P
- Graph Analysis Toolbox
- Pattern Discovery Toolbox
- first use of PB system

Answers to Moderator Questions

- What will be SGI's PFLOPS architecture?
 - Ultraviolet: GAM, efficient access, extreme synch, hetero processing
 - mass market content: processors, memory, network physical layer, OS
- How will we increase apps scalability?
 - petascale: efficient memory access, extreme synchronization
 - serial MATLAB: → parallel Star-P
- Size of HPC market? How can HPC customers help?
 - Much bigger if artificial barriers (*e.g.*, MPI/Fortran/C required) removed
 - Make large purchases an honest business
- Closest competitors?
 - Overall: today's, with cluster vendors becoming more prominent
 - Knowledge discovery: ???
- What will be SGI's HPC successes in 2009?
 - Establishment of bona fide knowledge discovery market
 - Connection with much larger market of users via Star-P

S&G is shifting the focus of its system architecture, to accelerating the discovery of patterns in immense data.

sggi[®]

Agenda

- Relevance of pattern discovery outside national security
- System requirements for pattern discovery
- Answers to moderator questions